

CINECA resources for scientific computing

Elda Rossi, Maurizio Cremonesi
CINECA - HPC and Innovation Unit
Italy



19 Giugno 2014

Agenda

- CINECA 2.0
- HPC computing infrastructure
- User environment @cineca
- Data analytics
- Questions and discussion



CINECA 2.0



CINECA
is a non profit Consortium,
made up of **69 Italian
universities***, and **3
Institutions** (CNR, OGS and
MIUR).

CINECA is now the largest Italian computing centre, one of the most important worldwide. .
The **High Performance Systems department (SCAI: SuperComputing Applications and Innovation)** offers support to scientific and technological research through supercomputing and its applications.

Data Storage



ICT Systems



Graphics



HPC Scientific



HPC Technical

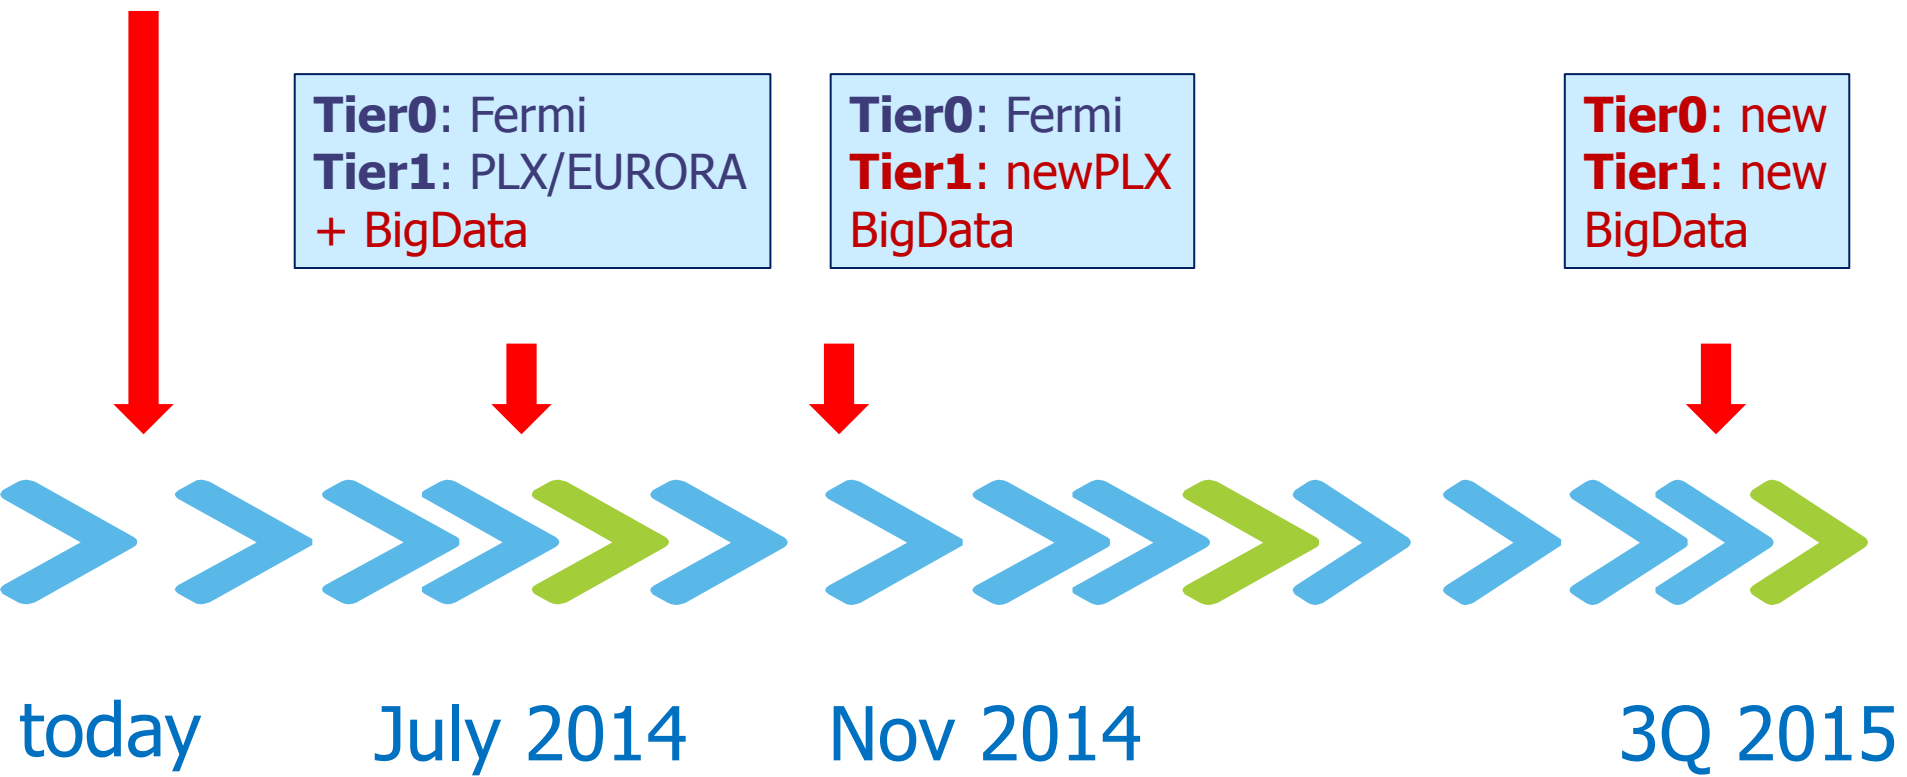


Tier0: Fermi
Tier1: PLX/EURORA

Tier0: Fermi
Tier1: PLX/EURORA
+ BigData

Tier0: Fermi
Tier1: newPLX
BigData

Tier0: new
Tier1: new
BigData



Name: Fermi
Architecture: BlueGene/Q (10 racks)
Processor type: IBM PowerA2 @1.6 GHz
Computing Nodes: 10.240
Each node: 16 cores and 16GB of RAM
Computing Cores: 163.840
RAM: 1GByte / core (163 TByte total)
Internal Network: 5D Torus
Disk Space: 2PByte of scratch space
Peak Performance: 2PFlop/s
Power Consumption: 820 kWatts
N. 12 in Top 500 rank (June 2013)
National and PRACE Tier-0 calls

High-end system, only
for extremely scalable
applications



Prototype system,
equipped with brand
new accelerators

Architecture: Hybrid cluster by EUROTECH

Processor type: Intel Xeon Sandy Bridge 3.1GHz

Computing Nodes: 64

Each node: 16 cores, 16GB/32 of RAM + 2 accelerators

Computing Cores: 1.024

RAM: 1GB/core

Accelerators: 64 NVIDIA Tesla K20 +
64 Intel Xeon-Phi 5120D (MIC)

Internal Network: Infiniband & Custom

Peak performance: 110 TFlops

Power consumption: 30 kWatts

N. 1 in Green 500 rank (June 2013)

National and PRACE PrepAccess calls



Name: PLX

Architecture: IBM Hybrid Cluster

Processor type: Intel Xeon Westmere @ 2.4 GHz

Computing Nodes: 274

Each node: 12 cores, 48GB of RAM, 2 GPUs

Computing Cores: 3.288

RAM: 14TByte

Internal Network: Infiniband 4xQDR switches (40 Gb/s)

Accelerators: 548 GPUs:

Peak Performance: 32 TFlops

565 TFlops SP GPUs

283 TFlops DP GPUs

X86 based system
for production of
medium scalability
applications



National and PRACE Tier-1 calls

(July 2014)

- Selected through a “call for tenders” with 5 participants
- IBM won the contract (to be officially accepted on June the 30th)
- Small-medium cluster (80 nodes) specially targeted to large dimension data analysis and storage, data analytics, pre-post processing applications (also web based) , scientific visualization, ...
- In general “less traditional” HPC applications (not only number crunching)
-

TAPE


- 12PB → 16PB
- New hw: 10 drives should guarantee 2.5GB/s throughput

DISKs

- 5PB distributed storage (GPFS) to be used across different platforms.
- Servers for Tiering and data migration

COMPUTE

- 70 nodes, 20 cores/each NeXtScale
- Intel Xeon E5-2600 v2 "Ivy Bridge"
- Mem: 128-512 GB/node
- 5 nodes BigInsight
- 40TB SSD disk



- New services to be defined on this system, taking advance from its peculiarities:
 - Low parallelism (less cores with respect to other systems, more cores/node)
 - Memory intensive (more memory/core and /node)
 - I/O intensive (SSD disk available)
 - DB based (a lot of storage space)
- New application environments:
 - Bioinformatics
 - Data analysis
 - Engineerings
 - Quantum Chemistry
- General services
 - Remote visualisation
 - Web access to HPC
 -



Name: ???

Architecture: Hybrid Cluster

Processor type: Intel Xeon Ivy Bridge (2.4 – 3GHz)

Computing Nodes: > 500

Computing Cores/cpu: 8 – 12 or more

RAM/core: >2 GB (section with more RAM)

Internal Network: Infiniband QDR (or FDR)

Accelerators: >500 GPUs Nvidia K40:

Peak Performance: 1 PFlops (>250 TF only cpu)

Cooling: liquid/ free cooling

Energy consuming: < 400 KW

- Fermi, at present our tier0 system, reaches the normal end
- It will be substituted with another system of comparable performance to fulfil the commitments at Italian and European level (order of magnitude 50PFlops -or- 50M€)
- BG/Q architecture is no more in the development plans of IBM, the actual technology has not yet been identified

How to get HPC resources

- Peer reviewed projects:
you can submit a project that will be reviewed. If you win you will get the needed resources for free
 - Local (Lombardia region) → LISA
 - National → ISCRA
 - Europe → PRACE
- No selection:
some Institutions got an amount of resources to be distributed among the research staff

Peer reviewed selection



LISA:

<http://www.hpc.cineca.it/services/lisa>

ISCRA:

<http://www.hpc.cineca.it/services/iscra>

PRACE:

<http://www.prace-ri.eu/Call-Announcements>

Basic rules (1)

The use of HPC resources is based on **projects** (or Account_no)

Each project is defined on the basis of:

- Budget (how many standard hours)
- Validity (from – to date)
- PI (Principal Investigator)
- Collaborators
- Host(s): one or more systems where the budget can be used

The project is active as far as the budget is not exhausted and in validity period.

The PI and Collaborators of a project can concurrently use the budget for batch processing.

Basic rules (1)

The **PI** can add collaborators to the project

Each user gets a **permanent username** (we call them "personal username") that enable him/her to access the HPC system/s

Access is normally done via ssh protocol giving the username/passwd information

All HPC systems share (if applicable) the same username/passwd information (if you change the passwd on one system, the change is spread to all the others)

The username is strictly **personal**. Please do not share it with collaborators or students. There is no problem in defining usernames for a lot of your collaborators or students.

Your **username** can be associated or not to projects, either as PI or Collaborator.

The username enables you to access the system and to work on it in **interactive mode** for compilations, editing, data movement, ... (please note: only a max of 10 min cpu time is allowed in this mode)

The real “hard production” has to be done using **batch mode**. You need a *valid project* in order to access this mode.

- Batch mode

The computing servers are used by many users all the time but:

- ✓ each user would like he/she to be the only user of the system
- ✓ or at least that others do not interfere with his/her jobs

A way for automatically realizing this is by using a *batch job management system*

The batch manager:

- looks at the users' jobs needs
- controls the available resources
- assign resources to each job
- in case put requests in a waiting batch queue

- Batch mode

The batch system needs the following infos per each job:

- ✓ which resources (nodes, cores, memory)
- ✓ for how much time

But the system administrator needs also to know who is paying for the job.

So the user must bundle his/her job with all these information.

- Write your script using an available editor

```
$ vi script
```

- The script file has 2 sections:
 - commands for the scheduler
(resources + Account_no)
 - commands for the system
(unix commands)

- Submit the script to the scheduler

```
$ qsub script
```

- Wait ... and check

```
$ qstat
```

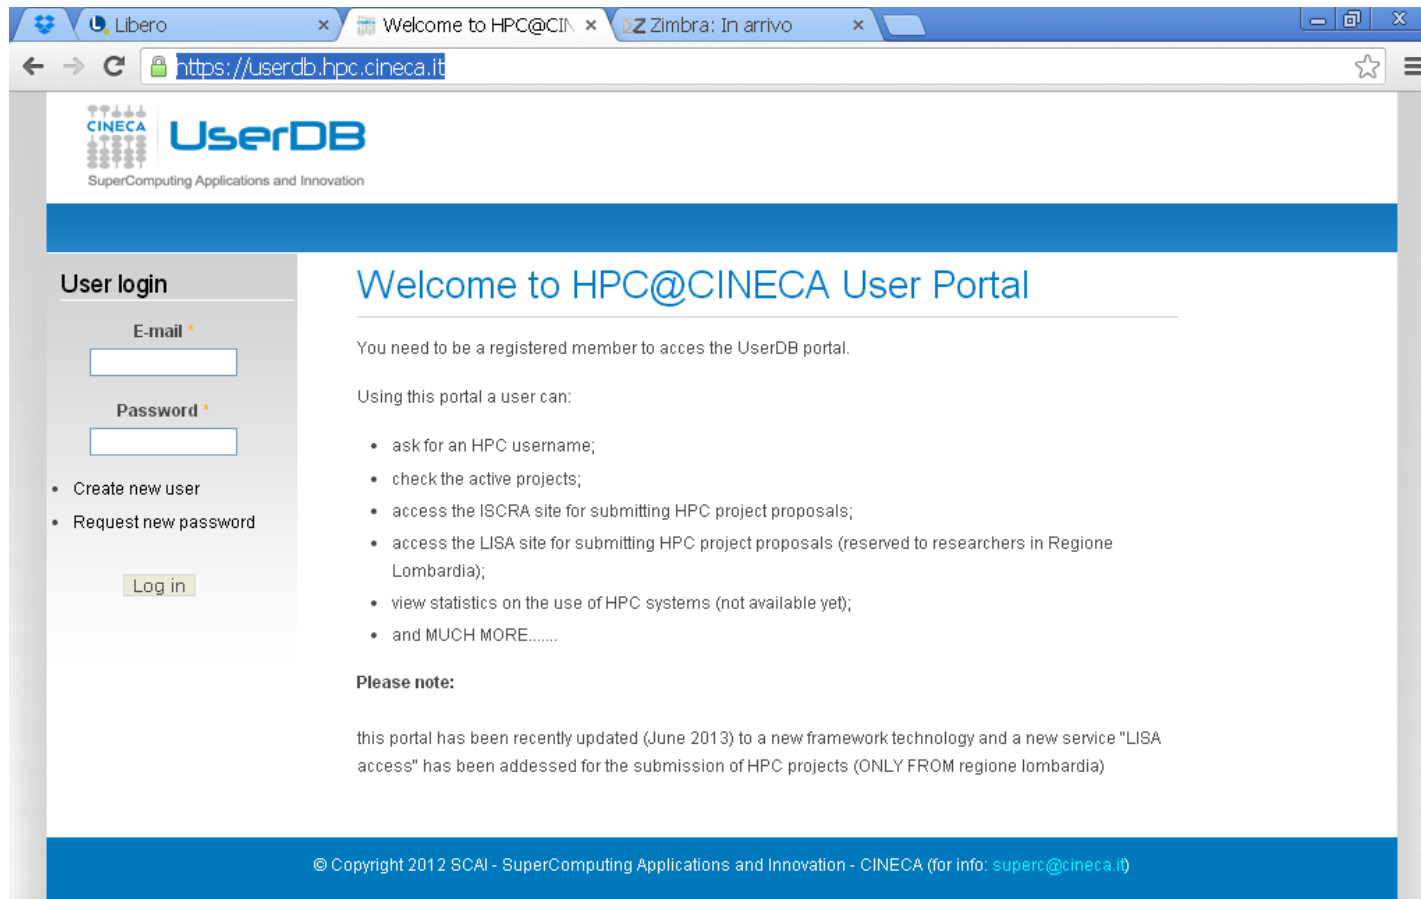
- The job completes: you can get final results

```
$ ls -l
```

```
#!/bin/bash
#PBS -l walltime=1:00:00
#PBS -l select=1:ncpus=4:mpiprocs=1
#PBS -o job.out
#PBS -e job.err
#PBS -q parallel
#PBS -A <my_account>

cd $PBS_O_WORKDIR
export TMPDIR =
module load stata
stata-mp -b do "input"
```

1. Register yourself to our UserDB (<https://userdb.hpc.cineca.it/>)
 - ... just once in your life!
 - Be ready to enter your CF (if italians) and your personal info
2. Get the association to a project
 - Ask the PI of a valid project to appoint you as a collaborator
 - or take part to a peer review selection, ...
3. Finalise your access request (in UserDB portal)
 - ... just once in your life!
 - In the UserDB environment, follow the "HPC access" link and accomplish the given instructions
 - Be ready to upload your ID card (or passport) - ... both sides!
 - In few hours an automatic mail will inform you about the access granting
 - In half a day your username will be available on the HPC system and/or the connection with the project(s) activated.



The screenshot shows a web browser window with the URL <https://userdb.hpc.cineca.it>. The page features the CINECA logo and the text "UserDB SuperComputing Applications and Innovation". On the left, there is a "User login" section with input fields for "E-mail" and "Password", and links for "Create new user" and "Request new password". A "Log in" button is also present. The main content area displays the heading "Welcome to HPC@CINECA User Portal" and a message stating that users need to be registered. A list of services available through the portal is provided, including requesting HPC usernames, checking active projects, and accessing proposal submission sites. A "Please note" section mentions a recent update in June 2013. The footer contains copyright information for SCAI and CINECA, along with the contact email superc@cinca.it.

<https://userdb.hpc.cineca.it/>

<http://www.hpc.cineca.it>

<http://www.hpc.cineca.it/content/hpc-user-guide>

The screenshot shows a web browser displaying the HPC User Guide page. The browser's address bar shows the URL www.hpc.cineca.it/content/hpc-user-guide. The page header includes the CINECA logo and the SCAI (SuperComputing Applications and Innovation) logo. A navigation menu is visible with categories: ABOUT US, RESOURCES, SERVICES, FOR USERS (selected), TRAINING, and PROJECTS. A search bar is located in the top right corner. The main content area features a status bar with three columns: FERMI status, PLX status, and Eurora status, each with three colored circles (green, yellow, red). Below this, the breadcrumb trail reads 'Home > For users > Documentation'. The main heading is 'HPC User Guide'. Under the heading, there is a 'Content:' section with a list of links: Introduction, General Info, and System Specific. At the bottom of the page, a subtitle reads 'A User Guide for HPC systems in CINECA (last release: january 2014)'. A left sidebar contains a 'For users' section with a list of links: UserDB, Getting started, Get in touch, Help desk, and Documentation (selected). Under 'Documentation', there is a link for 'HPC User Guide' and a list of sub-topics: Introduction, General Info, System Specific, FAQ, Other Documents, Services, One-page manuals, and Documentation at a glance.

- Send a mail to superc@cineca.it
 - Always an on-call operator during working hours
 - You can also ask for domain-driven questions: the operator will find the right consultant for you

The screenshot shows the SCAI (SuperComputing Applications and Innovation) Help Desk website. The browser address bar shows the URL www.hpc.cineca.it/content/help-desk. The page features a navigation menu with categories: ABOUT US, RESOURCES, SERVICES, FOR USERS (selected), TRAINING, and PROJECTS. A search bar is located in the top right corner. The main content area displays the 'Help desk' page, which includes a 'For users' sidebar with links to 'My portal', 'Getting started', 'Get in touch', 'Help desk' (highlighted), and 'Documentation'. The main content area shows the 'Help desk' title and a list of consultants: Isabella Baccarelli, Mirko Cestari, Fabrizio Cinquini, Francesco Falciano, Silvia Giuliani, Alessandro Grottesi, Giusy Muscianisi, Nicola Spallanzani, and Elda Rossi. Above the list, there are status indicators for 'FERMI status' and 'PLX status', each represented by three colored circles (green, yellow, red).



For users

- » My portal
- » Getting started
- » **Get in touch**
- » Help desk
- » Documentation

Help desk



Alessandro Grottesi

Center news

30/10/2013

Crush of some running jobs on FERMI

28/10/2013

Fermi Scratch nearly full

FERMI status



PLX status



[Home](#) > [For users](#)

Get in touch

How to get Center's announcements (HPC-news)

We manage a mailing list (HPC-news) for posting announcements, scheduled downs, software updates, any problems and so on, about our HPC computing resources. It is advisable for HPC users to be included in that list!

You can subscribe (or unsubscribe) to HPC-news by sending an email from the address you want to subscribe. You can consult the archive browsing the archive web site.

To subscribe to HPC-new:

send a mail to listserv@list.cineca.it
 in the body --> "subscribe hpc-news"
 in the subject --> any string...

To unsubscribe to HPC-new:

send a mail to listserv@list.cineca.it
 in the body --> "unsubscribe hpc-news"
 in the subject --> any string...

To consult the archive:

for the recent archive go to → [Center News](#)
 for the full archive go to → <http://list.cineca.it/archives/hpc-news.html>

Access to the system

1. Interactive access: ssh client
2. Access via interface:
 - Web-based via WebCompute
 - RCM: remote connection manager
3. Data Transfer: sftp client

Ssh (Secure Clients)

Shell for Linux users (scp, ssh)

For windows users:

- Putty (ssh)
- TECTIA client (ssh)
- Winscp for windows (sftp)

PLX login: *ssh login.plx.cineca.it*

```
Last login: Wed Oct 30 08:35:17 2013 from 131.175.80.185
```

```
*****  
*  
* Welcome to PLX DataPlex Cluster @ CINECA - RedHat EL 5.6! *  
*  
* Qlogic QDR (40Gb/s) Infiniband high-performance network *  
*  
* 274 Compute node *  
* - 2 esa-core Intel(R) Xeon(R) CPU E5645 @2.40GHz per Compute node *  
* - 48 GB RAM per Compute node *  
* - 2 Nvidia Tesla M2070 GPU per Compute node *  
* 8 Fat node *  
* - 2 quad-core Intel(R) Xeon(R) CPU X5570 @2.93GHz per Fat node *  
* - 128 GB RAM per Fat node *  
* 3352 Total cores *  
*  
* 6 Remote Visualization Login *  
* 2 Nvidia QuadroPlex 2200 S4 *  
*  
* PBSpro 10.4 batch scheduler *  
*  
* http://www.hpc.cineca.it/content/ibm-plx-gpu-user-guide-0 *  
* for a guide on PLX *  
*  
* mailto:superc@cinca.it for support *  
*  
*****  
[rponzini@node342 ~]$ █
```

Some applications are pre-installed, using the
"module" environment

```
$ module available  
$ module load stata  
$ module help stata
```

```
Module Specific Help for /cineca/prod/modulefiles/base/applications/stata/10:
```

```
...  
Stata is a complete, integrated statistical package that provides ...  
...  
In the following you can find a batch job script example for stata-mp:
```

```
=====  
#!/bin/bash  
#PBS -l walltime=1:00:00  
#PBS -l select=1:ncpus=4:mpiprocs=1  
#PBS -o job.out  
#PBS -e job.err  
#PBS -q parallel  
#PBS -A <my_account>
```

```
cd $PBS_O_WORKDIR  
export TMPDIR =  
module load stata  
stata-mp -b do "input"
```

```
...
```

\$ module available

```

----- /cineca/prod/modulefiles/profiles -----
profile/advanced          profile/base(default)          profile/engineering
.....
----- /cineca/prod/modulefiles/base/libraries -----PETSc/3.0.0--openmpi--1.3.3--
gnu--4.1.2
PETSc/3.0.0--openmpi--1.3.3--intel--11.1--binary(default)

----- /cineca/prod/modulefiles/base/compilers -----
IntelMPI/4.0--binary          gnu/4.1.2(default)
gnu/4.5.2                    intel/11.1--binary(default)
openmpi/1.3.3--gnu--4.1.2    openmpi/1.3.3--intel--11.1--binary

----- /cineca/prod/modulefiles/base/applications -----
R/2.10.1                    gromacs/4.6.1
abinit/6.12.3                meep/1.1.1(default)
adf/2010.02b(default)        molcas/7.6
amber/11(default)           mopac/2009(default)
amber/12                    namd/2.8(default)
cp2k/2.3(default)           pyfrag/2007.02(default)
desmond/3.0.3                stata/10
....

```

- Username
 - Personal (erossi00, rponzini, mcremone)
 - For access and interactive use
 - It will be closed one year after any valid projects
- Project = Account_no
 - Budget per the batch mode
 - validity, budget, hosts, PI, collaborators
 - "saldo -b" gives you a list of your projects

```
$ saldo -b
```

```
-----
account      start      end        total      localCluster  totConsumed  totConsumed
              (local h)  Consumed(local h)  (local h)  %
-----
try11_test   20110301  20111201  10000      2000          5000         50.0
cin_staff    20110323  20200323  200000000  479621        8933910      4.5
```

```
$ saldo -ba try11_test
```

```
erossi00 ffalcian mcestari fvitale0 sgiulian
```

```
-----
account      start      end        total      localCluster  totConsumed  totConsumed
              (local h)  Consumed(local h)  (local h)  %
-----
try11_test   20110301  20111201  10000      2000          5000         50.0
```

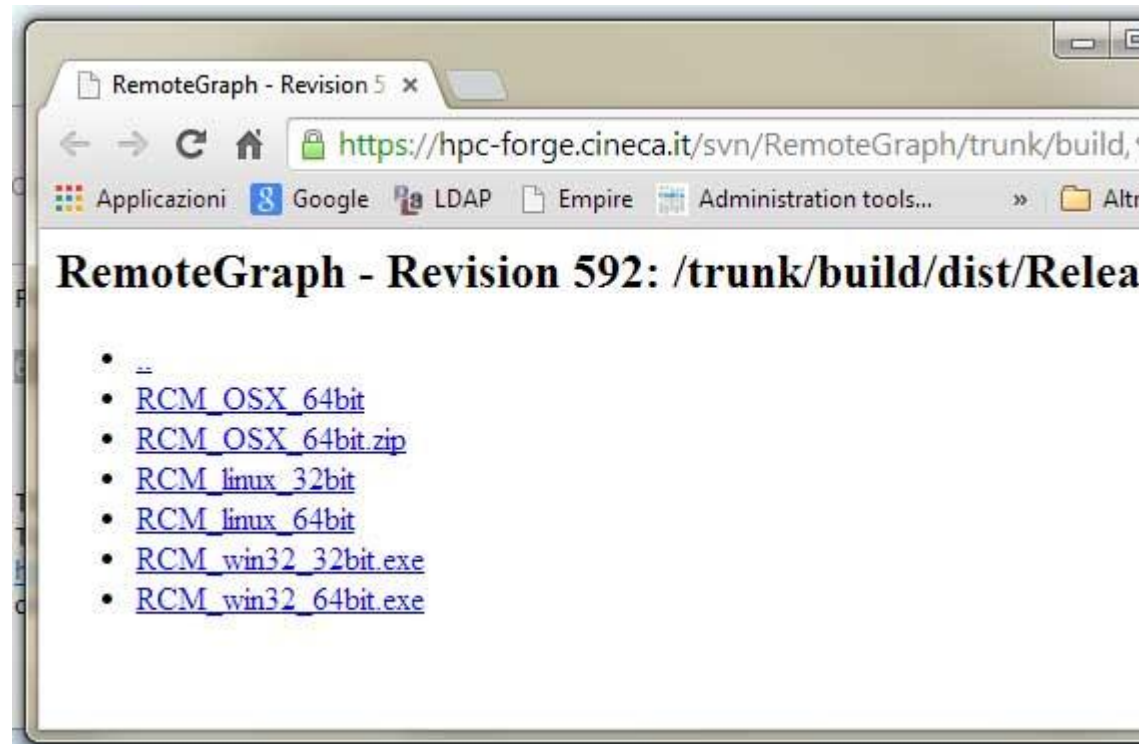

GUI access and remote visualization

- Two different tools, both of them make use of accelerators for remote visualization
 - **RCM: Remote Connection Manager**
 - ***Webcompute***

Remote Connection Manager

<http://www.hpc.cineca.it/content/remote-visualization>

- Dept. web site
www.hpc.cineca.it
- Follow the link
Services → RemoteVisualiz
→ download
- Download the correct client
- Execute it



Remote Connection Manager

The image displays the Remote Connection Manager (RCM) interface, which is used for managing remote connections to a cluster. It is divided into three main sections:

- RCM Login:** A window for logging into the system. It features the CINECA logo and the text "REMOTE CONNECTION MANAGER version: 1.1.365". The login fields are: Sessions (rponzini@login.plx.cineca.it), Host (login.plx.cineca.it), User (rponzini), and Password (masked with asterisks). A "LOGIN" button is located at the bottom.
- Remote Connection Manager 1.1.365 - CINECA:** A window showing the status of active connections. It includes a table with columns for CREATED, DISPLAY, NODE, STATE, TIMELEFT, USERNAME, and WALLTIME. The table contains one entry for a connection to node097. Below the table are buttons for "CONNECT", "KILL", "NEW DISPLAY", and "REFRESH". The status "Idle" is shown at the bottom left.
- TurboVNC:** A window showing a remote desktop view of a Linux system. The desktop environment is Konqueror. The file manager shows a directory containing several desktop files: paraview3.14.desktop, paraview3.98.desktop, paraview4.0.1.desktop, paraview_demo1.desktop, tecplot.desktop, UnigineGraph icTest.desktop, and Vaa3D.desktop. The status bar at the bottom indicates "7 Items - 7 Files (2.9 KB Total) - No Folders".

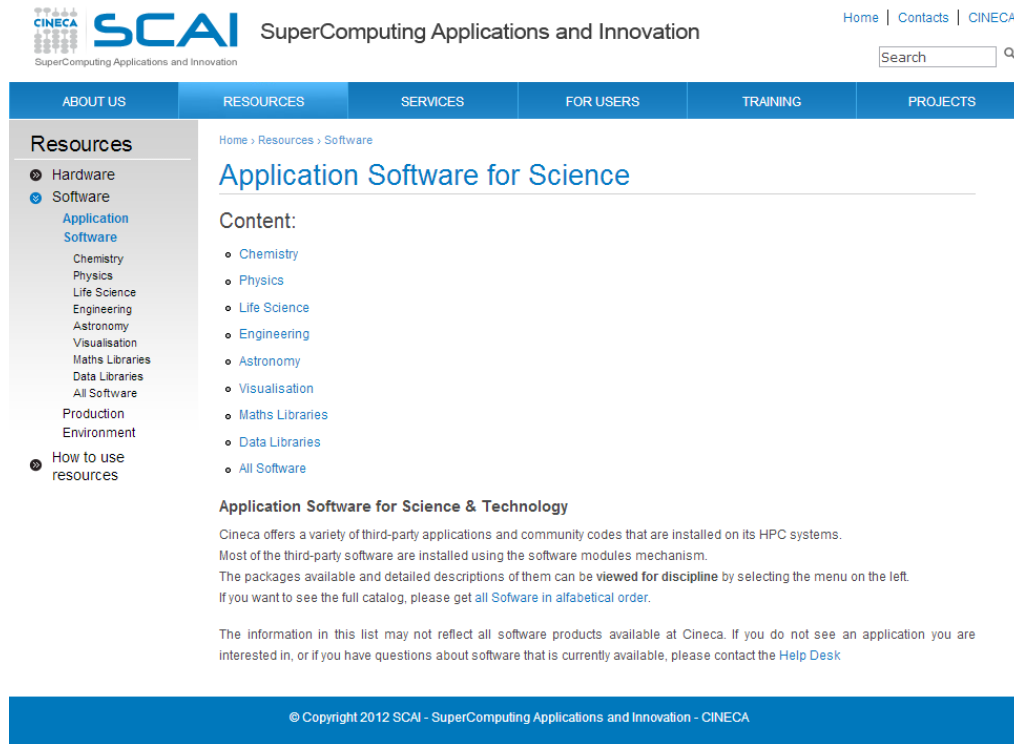
Red arrows indicate the flow of the process: from the login window to the connection manager window, and from the connection manager window to the remote desktop view.

- webcompute.cineca.it
- Download the correct client (only once)
- Select the service of interest (*xterm* o other applications)
- Select the project you want to account to and submit

The image shows a screenshot of the CINECA webcompute interface and a VNC viewer window. The web interface is at the top left, showing the URL <https://webcompute.cineca.it/engineframe/runyourjob/it.cineca.runyourjob>. The page has a navigation menu with 'Home', 'My Sessions', 'My Data', and 'My Jobs'. Under 'HPC Services', there is a list of services: XTerm, Pointwise, Pointwise 16, and ParaView. The main content area says 'Welcome to CINECA's Computing Services' and provides instructions for downloading VNC clients for Windows, Linux, and Mac OS X. A red arrow points from the 'XTerm' service in the web interface to a VNC viewer window on the right. The VNC viewer window shows a terminal window with the prompt `[erossi00@node098 ~]$`.

Pre-installed software

<http://www.hpc.cineca.it/content/application-software-science>



CINECA **SCAI** SuperComputing Applications and Innovation

Home | Contacts | CINECA

Search

ABOUT US | RESOURCES | SERVICES | FOR USERS | TRAINING | PROJECTS

Resources

- Hardware
- Software
 - Application Software
 - Chemistry
 - Physics
 - Life Science
 - Engineering
 - Astronomy
 - Visualisation
 - Maths Libraries
 - Data Libraries
 - All Software
 - Production
 - Environment
- How to use resources

Home > Resources > Software

Application Software for Science

Content:

- Chemistry
- Physics
- Life Science
- Engineering
- Astronomy
- Visualisation
- Maths Libraries
- Data Libraries
- All Software

Application Software for Science & Technology

Cineca offers a variety of third-party applications and community codes that are installed on its HPC systems. Most of the third-party software are installed using the software modules mechanism. The packages available and detailed descriptions of them can be viewed for discipline by selecting the menu on the left. If you want to see the full catalog, please get all Software in alphabetical order.

The information in this list may not reflect all software products available at Cineca. If you do not see an application you are interested in, or if you have questions about software that is currently available, please contact the Help Desk

© Copyright 2012 SCAI - SuperComputing Applications and Innovation - CINECA

